

## CLAIMS

What is claimed as new and desired to be protected by Letters Patent of the United States is:

1. A method for processing a host command in a storage system including a plurality of storage elements coupled to a plurality of storage controllers, the method comprising the steps of:

at a first one of said plurality of storage controllers,

receiving a host command directed to a volume of said storage system;

determining a target storage element of said storage system corresponding to said volume;

determining a target storage controller from said plurality of storage controllers corresponding to said target storage element;

if said target storage controller is not said first one of said plurality of storage controllers, forwarding the host command to said target storage controller;

at the target storage controller,

receiving the host command forwarded by the first one of said plurality of controllers;

executing said host command;

forwarding, to said first one of said plurality of storage controllers, an execution status associated with said host command; and

at said first one of said plurality of storage controllers,

receiving said execution status; and

forwarding said execution status to said host.

2. The method of claim 1, wherein said step of determining a target storage controller comprises searching a mapping table.

3. The method of claim 2, further comprising the step of:

before receiving a host command at said first one of said plurality of storage controllers, creating said mapping table to reflect a mapping between volumes and controllers of said storage system.

4. A method for processing a host command in a storage system including a plurality of storage elements coupled to a plurality of storage controllers, the method comprising the steps of:

at a first one of said plurality of storage controllers,

receiving a host command directed to a virtual volume of said storage system;

determining a plurality of logical volumes of said storage system corresponding to said virtual volume;

determining a plurality of target storage controllers from said plurality of storage controllers corresponding to said plurality of logical volumes;

sending a component command to each of said plurality of target storage controllers, the component command for a particular one of said plurality of target storage controllers corresponding to at least a portion of the host command and relating to the logical volume associated with the particular target controller;

at each of said plurality of the target storage controllers,

receiving the component command sent by the first one of said plurality of controllers;

executing said component command;

forwarding, to the first one of said plurality of storage controller, an execution status; and

at said first one of said plurality of storage controllers,

receiving said execution status from each one of said plurality of target storage controllers;

determining an aggregate host command execution status from said received execution status; and

forwarding said aggregate host command status to said host.

5. The method of claim 4, wherein said step of determining a plurality of target storage controllers comprises searching a mapping table.

6. The method of claim 5, further comprising the step of:

before receiving a host command at said first one of said plurality of storage controllers, creating said mapping table to reflect a mapping between virtual volumes and logical volumes of said storage system.

7. A storage system, comprising:

a plurality of storage controllers, each one of said plurality of storage controllers comprising:

at least one host port for communicating with a plurality of hosts; and

at least one storage element port for communicating with a plurality of storage elements;

an interconnect connected to each of said plurality of storage controllers; and

at least one storage element, each storage element being coupled to a respective storage element port;

wherein said plurality of storage controllers further comprise means for processing a host command received on any host port targeted to a storage volume on any one of said plurality of storage elements.

8. The storage system of claim 7, further comprising:

a configuration controller for setting up mappings between storage volumes and storage elements.

9. A storage system, comprising:

a plurality of storage controllers, each one of said plurality of storage controllers comprising:

at least one host port for communicating with a plurality of hosts; and

at least one storage element ports for communicating with a plurality of storage elements;

an interconnect connected to each of said plurality of storage controllers; and

at least one storage element, each storage element being coupled to a respective storage element port;

wherein said plurality of storage controllers further comprise means for processing a host command received on any host port targeted to a virtual volume on said plurality of storage elements.

10. The storage system of claim 9, further comprising:

a configuration controller for setting up mappings between virtual volumes and logical volumes.

11. A scalable storage controller, comprising:

a high speed interconnect;

a plurality modules, each of said plurality of modules comprising:

a cache memory;

an interconnect port, coupled to said high speed interconnect;

a storage port, for coupling to a storage device;

a host port, for coupling to a host; and

a processing element, coupled to said cache memory, said interconnect port, said storage port, and said host port.

12. The scalable storage controller of claim 11, wherein said high speed interconnect is a serial interconnect.

13. The scalable storage controller of claim 11, wherein said high speed interconnect is a fibre channel loop.

14. The scalable storage controller of claim 11, wherein said high speed interconnect is a switched serial network.

15. A scalable storage controller, comprising:

a primary interconnect; and

a plurality of modules, each of said plurality of modules comprising:

a cache memory;

an interconnect port, coupled to said primary fibre channel interconnect;

a storage port, for coupling to a storage device;

a host port, for coupling to a host; and

a processing element, coupled to said cache memory, said interconnect port, said storage port, and said host port;

wherein the processing element of each of said modules are configured to process a host command for accessing a virtual volume by causing each module associated with said virtual volume to access its respective storage device.

16. The scalable storage controller of claim 15, further comprising:

a redundant interconnect;

wherein each of said modules further comprises another interconnect port coupled to said redundant interconnect.

17. In a storage controller comprising a plurality of modules each capable of receiving host commands, a method of conflict detection, comprising:

receiving, at a first module, a first access request;

receiving, at a second module, a second access request;

at each of said first and second modules,

determining whether said received access request requires coordination;

if said received access request is determined to require coordination, executing a coordination request, and

if said coordination request is not granted,

indicating in a conflict table that said received access request is in conflict; and

holding said access request;

if said coordination request is granted,

executing said received access request;

searching said conflict table to find an entry in conflict with said received access request;

if said search is successful, notifying the module associated with said found entry; and

if said search is unsuccessful, releasing said coordination request.

18. The method of claim 17, wherein said step of executing a coordination request comprises:

searching a coordination data structure to find a data structure having an address range



encompassing an address range associated said received access request;

if the data structure is found, attempting to reserve said data structure;

if said attempt to reserve is not successful,

indicating the said coordination request is not granted;

if said attempt to reserve is successful,

indicating that said coordination request is granted;

if said data structure is not found,

allocating said data structure and indicating said coordination request is granted.